4.5 A 100mW 9.6Gb/s Transceiver in 90nm CMOS for Next-Generation Memory Interfaces

Edoardo Prete, Dirk Scheideler, Anthony Sanders

Infineon, Munich, Germany

To address the bandwidth requirements of next-generation memory systems, a distributed architecture has been adopted with an advanced memory buffer chip (AMB) on every dual inline memory module (DIMM). The AMB acts as an interface hub connecting the memory controller to each DIMM in a daisy chain, and transfers memory requests to the local DRAM via a standard DDR interface. The hub-to-hub connection is a wide high-speed interface currently standardized at 4.8Gb/s, and will be pushed to 9.6Gb/s in the next generation. Although many previous memories have been demonstrated at this speed, the power, reliability and cost of such a system have forced innovation in order to provide the necessary leading-edge solution.

The cost of system interconnect demands cheap materials and connectors, and requires equalization at the transmitter and good jitter performance. Due to the poor noise performance of low-cost clock buffers, a forwarded clock architecture is necessary above 4.8Gb/s where, parallel to the wide high-speed data bus, a half baud-rate clock is transmitted with inherently the same jitter components as the data.

We will present in this paper a high-performance low-power transceiver architecture with clock distribution targeted to meet the requirements for next-generation memory systems. Implemented in 90nm CMOS technology, 10Gb/s operation is achieved with 100mW/channel power consumption and a 1.0V supply. Figure 4.5.1 shows a block diagram of the transceiver together with the clock receiver and distribution and how the forwarded clock is propagated through the daisy-chained system.

The CDR architecture is based on the forwarded clock and, by means of a digitally-programmable PLL [1], optimally tracks the relatively high-frequency wander of the data, up to 150MHz, by using the common correlated phase information of the forwarded clock. The CDR circuit tracks mid-band wander present only in the data, and rejects any high-frequency jitter present on either the clock or data. The 150MHz tracking bandwidth is a compromise between tracking jitter induced through package resonances at the transmitter, and filtering any correlated jitter components, which have become out of phase because of clock and data skew.

The detailed CDR block diagram is shown in Fig. 4.5.2, where a PLL and a phase-recovery loop can be identified. The PLL locks to the forwarded clock and provides a low-pass phase filter function. It uses a half-rate 4-stage CML ring oscillator that generates 8 evenly-spaced phases. Four linear-phase detectors compare the forwarded clock phase to each of 4 evenly-spaced phases from the VCO, the outputs of which are then weighted and added as currents in the loop filter. By adjusting this weighting, the locking phase of the PLL can be steered without bound over 360°. The PLL is additionally provided with a frequency-acquisition loop using divided-down clocks to detect cycle slips and acting via an additional charge pump on the loop-filter capacitor. An Alexander phase detector in conjunction with a digital filter elaborates the phase error information of the data and steers the weights of the phase detectors, in order to align the output clock. The effective bandwidth of this loop is set to be relatively low as only low-frequency wander components, less than 500kHz, are present due to voltage and temperature variations.

The charge pump and loop filter use a differential architecture to improve noise rejection and the VCO incorporates interpolating CML stages with differential tuning. The loop filter is driven by the digitally-controlled bias differential XOR phase detectors, giving a total of 1/32 unit-interval phase resolution.

To confirm correct operation of the CDR circuit, the tracking behavior is measured in the laboratory using a dual pattern generator with voltage-controlled delay lines. The forwarded clock and data presented to the chip can, using this setup, be sinusoidal phase modulated at different frequencies and jitter amplitudes. For 9.6Gb/s operation and realistic received signal amplitude, $<100mV_{pp}$ differential, the modulation frequency is swept and the amplitude increased until a bit error rate (BER) of 10e-12 is confidently measured. By modulating only the data, Fig. 4.5.3, the low-pass tracking bandwidth of the CDR circuit can be confirmed as 500kHz, and demonstrates a jitter tolerance of 61ps for the receiver. Through correlated phase modulation of the clock and data, Fig. 4.5.3, the low-pass tracking bandwidth of the PLL can be confirmed as 150MHz with no peaking. At the point where the PLL no longer tracks, the effective jitter seen by the data receiver is then increasing and causes the roll-off in tolerable jitter. The measured jitter on the recovered clock is 21ps including 6.5ps (2 steps) of digital control noise.

Referring to Fig. 4.5.4, the ground-referenced transmitter features a 2-tap programmable de-emphasis and amplitude regulation. To achieve the output linearity and return loss a CML output stage was implemented, however, because of the ground termination, this had to be implemented using PMOS devices, which suffer from performance degradation compared to NMOS-based transmitters [2]. In order to achieve the switching speeds and reduce power, a full swing pre driver is implemented featuring a swing control to limit the output common-mode noise and duty-cycle distortion. The final 2-to-1 multiplexer uses a half-rate clock and the shift register is latch based to reduce power consumption and latency respectively. By implementing the driver using 32 individually switchable parallel stages, the amplitude of the taps is accurately programmed and controlled using a replica circuit.

The transmitter performance was confirmed in the lab by observing the output eye for a PRBS31 pattern on an oscilloscope triggered by the reference clock. By optimally programming the taps of the transmitter to account for the test board, a total output jitter of 19ps was measured, Fig. 4.5.5, and by over emphasizing the transmitter, the de-emphasis operation of the transmitter is confirmed.

The clock distribution network over 4mm was realized by a single fly-by pseudo transmission line constructed using lumped inductors. The network is represented in Fig. 4.5.6, and is divided into 7 segments each serving 2 transceivers and the next segment. Each segment consists of a differential metal line with additionally inserted series inductors, which compensate for the *RC* losses and lumped loads. Using such an approach the inherent propagation properties of a transmission line are utilized to distribute the clock with 60% less power than a traditional CML approach. A chip micrograph is shown in Fig. 4.5.7.

Acknowledgments:
We would like to thank M. Schreiner, M. Ebner, J. Blank, C. Andreotti & N. Unger for their contribution.

References:
[1] T. Toifl et al., "0.94ps-rms-Jitter 0.016mm² 2.5GHz Multi-Phase Generator PLL with 360° Digitally Programmable Phase Shift for 10Gb/s Serial Links," *IEEE ISSCC Dig. Tech. Papers*, pp. 410-412, Feb., 2005.
[2] R. Farjad-Rad et al., "A 0.4-µm CMOS 10-Gb/s 4-PAM Pre-Emphasis Serial Link Transmitter," *IEEE J. Solid-State Circuits*, vol. 34, pp. 580–585, May, 1999.

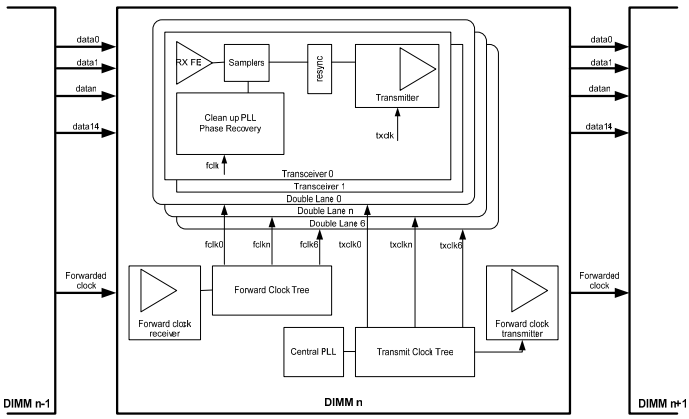

Figure 4.5.1: Overall block diagram.

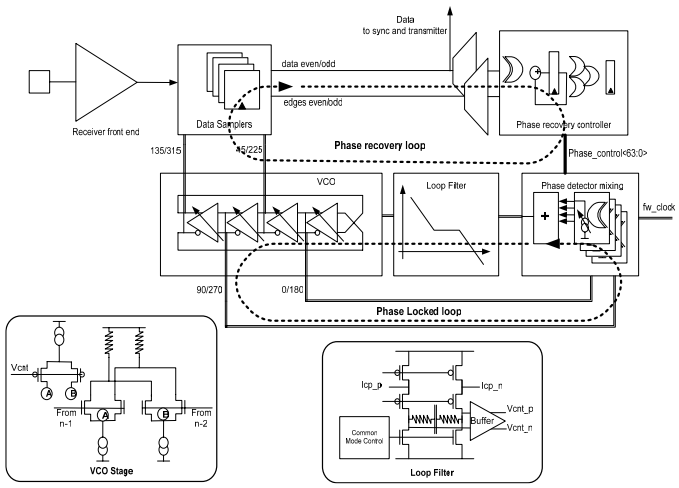

Figure 4.5.2: CDR block diagram.

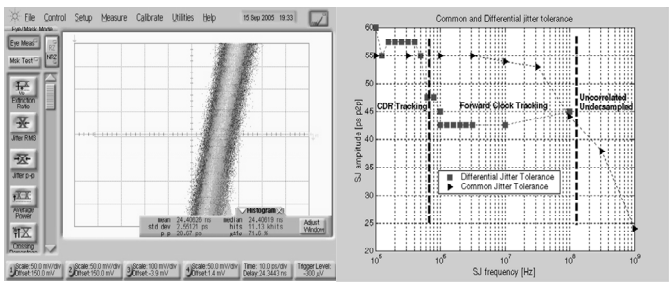

Recovered Clock Jitter

Jitter Tolerance

Figure 4.5.3: CDR performance.

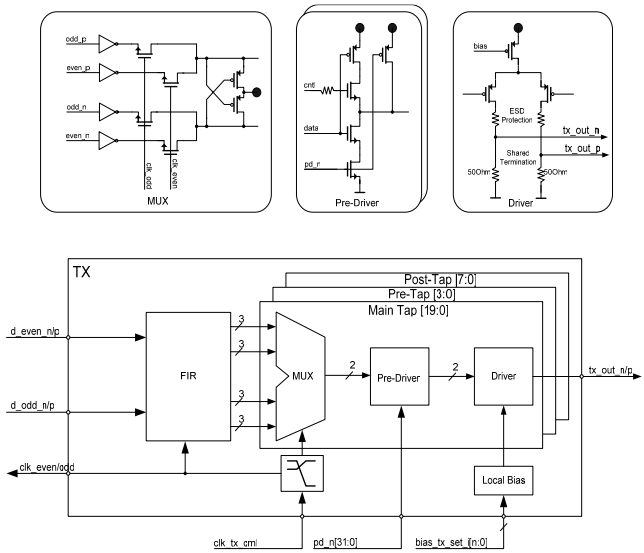

Figure 4.5.4: Transmitter block diagram.

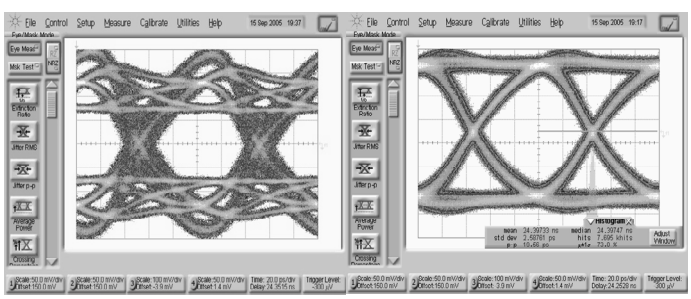

Full de-emphasis
10 Gbps over test board

Optimal de-emphasis
10 Gbps over test board

Figure 4.5.5: Transmit eye.

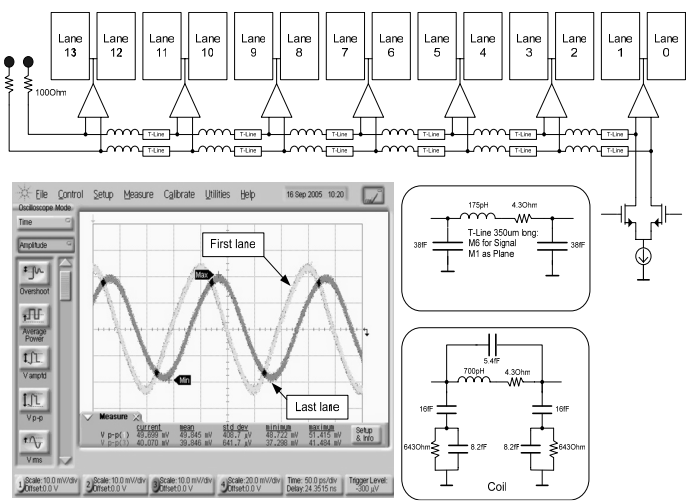

Figure 4.5.6: Clock tree.

Continued on Page 640

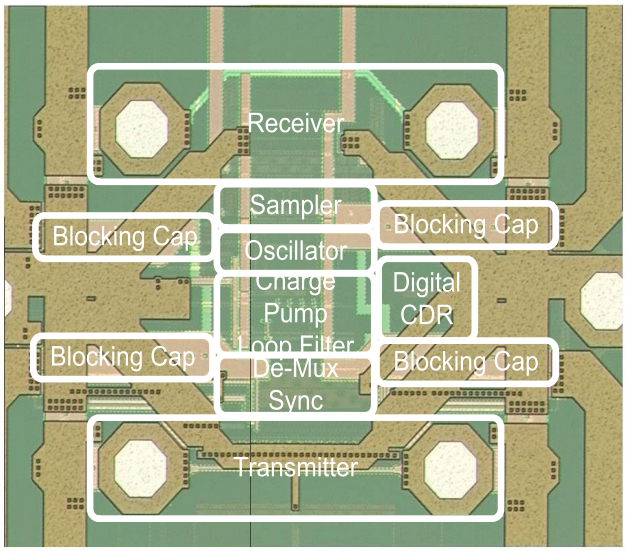

Figure 4.5.7: Chip micrograph.

